PATENT APPLICATION

DENTAL DATA MINING

Inventor(s):   ERIC E. KUO, a citizen of United States,
residing at 912 Beach Park Boulevard, #86
Foster City, California 94404;

PHILIPPE DE SMEDT, a citizen of Belgium,
residing at 1179 Arrowfield Way
San Ramon, California 94583;

CUONG VAN NGUYEN, a citizen of United States,
residing at 2227 McLaughlin Avenue, #2
San Jose, California 95122;

CHRISTOPHER W. OVERTON a citizen of United States,
residing at 1312 Fernside Boulevard
Alameda, California 94501.

Assignee: ALIGN TECHNOLOGY, INC.
881 Martin Avenue
Santa Clara, CA 95050-2903

Status: Large Entity

CORRESPONDENCE ADDRESS
Bao Q. Tran, Esq.
Align Technology, Inc.
881 Martin Avenue
San Jose, CA 95050
Tel.: (408) 470-1243
Fax: (408) 470-1024

ATTORNEY DOCKET: **DATE OF DEPOSIT:**   February 26, 2004

AT-000224 US   **EXPRESS MAIL NO.:**   **EL 995398942 US**

DENTAL DATA MINING

BACKGROUND

The present invention relates to computational orthodontics and dentistry.

In orthodontic treatment, a patient's teeth are moved from an initial to a final position using any of a variety of appliances. An appliance exerts force on the teeth by which one or more of them are moved or held in place, as appropriate to the stage of treatment.

SUMMARY

Systems and methods are disclosed providing a database comprising a compendium of at least one of patient treatment history; orthodontic therapies, orthodontic information and diagnostics; employing a data mining technique for interrogating said database for generating an output data stream, the output data stream correlating a patient malocclusion with an orthodontic treatment; and applying the output data stream to improve a dental appliance or a dental appliance usage.

The achieved outcome, if measured, is usually determined using a set of standard criteria such as the American Board of Orthodontics against which the final outcome is compared to, and is usually a set of idealized norms of what the ideal occlusion and bite relationship ought to be. Another method of determining outcome is to use a relative improvement index such as PAR, IOTN, and ICON to measure degrees of improvement as a result of treatment.

The present invention provides methods and apparatus for mining relationships in treatment outcome and use the mined data to enhance treatment plans or enhance appliance configurations in a process of repositioning teeth from an initial tooth arrangement to a final tooth arrangement. The invention can operate to define how repositioning is accomplished by a series of appliances or by a series of adjustments to appliances configured to reposition individual teeth incrementally. The invention can be applied advantageously to specify a series of appliances formed as polymeric shells having the tooth-receiving cavities, that is, shells of the kind described in the above-mentioned U.S. Application No. 09/169276, (Attorney Docket No. 018563-004800US-AT-00105US), filed Oct. 8, 1998.

A patient's teeth are repositioned from an initial tooth arrangement to a final tooth arrangement by making a series of incremental position adjustments using appliances specified in accordance with the invention. In one implementation, the invention is used to specify shapes for the above-mentioned polymeric shell appliances. The first appliance of a series will have a geometry selected to reposition the teeth from the initial tooth arrangement to a first intermediate arrangement. The appliance is intended to be worn until the first intermediate arrangement is approached or achieved, and then one or more additional (intermediate) appliances are successively placed on the teeth. The final appliance has a geometry selected to progressively reposition teeth from the last intermediate arrangement to a desired final tooth arrangement.

The invention specifies the appliances so that they apply an acceptable level of force, cause discomfort only within acceptable bounds, and achieve the desired increment of tooth repositioning in an acceptable period of time. The invention can be implemented

to interact with other parts of a computational orthodontic system, and in particular to interact with a path definition module that calculates the paths taken by teeth as they are repositioned during treatment.

In general, in one aspect, the invention provides methods and corresponding apparatus for segmenting an orthodontic treatment path into clinically appropriate substeps for repositioning the teeth of a patient. The methods include providing a digital finite element model of the shape and material of each of a sequence of appliances to be applied to a patient; providing a digital finite element model of the teeth and related mouth tissue of the patient; computing the actual effect of the appliances on the teeth by analyzing the finite elements models computationally; and evaluating the effect against clinical constraints. Advantageous implementations can include one or more of the following features. The appliances can be braces, including brackets and archwires, polymeric shells, including shells manufactured by stereo lithography, retainers, or other forms of orthodontic appliance. Implementations can include comparing the actual effect of the appliances with an intended effect of the appliances; and identifying an appliance as an unsatisfactory appliance if the actual effect of the appliance is more than a threshold different from the intended effect of the appliance and modifying a model of the unsatisfactory appliance according to the results of the comparison. The model and resulting appliance can be modified by modifying the shape of the unsatisfactory appliance, by adding a dimple, by adding material to cause an overcorrection of tooth position, by adding a ridge of material to increase stiffness, by adding a rim of material along a gumline to increase stiffness, by removing material to reduce stiffness, or by redefining the shape to be a shape defined by the complement of the difference between

the intended effect and the actual effect of the unsatisfactory appliance. The clinical constraints can include a maximum rate of displacement of a tooth, a maximum force on a tooth, and a desired end position of a tooth. The maximum force can be a linear force or a torsional force. The maximum rate of displacement can be a linear or an angular rate of displacement. The apparatus of the invention can be implemented as a system, or it can be implemented as a computer program product, tangibly stored on a computer-readable medium, having instructions operable to cause a computer to perform the steps of the method of the invention.

Among the advantages of the invention are one or more of the following. Appliances specified in accordance with the invention apply no more than orthodontically acceptable levels of force, cause no more than an acceptable amount of patient discomfort, and achieve the desired increment of tooth repositioning in an acceptable period of time. The invention can be used to augment a computational or manual process for defining tooth paths in orthodontic treatment by confirming that proposed paths can be achieved by the appliance under consideration and within user-selectable constraints of good orthodontic practice. Use of the invention to design aligners allows the designer (human or automated) to finely tune the performance of the aligners with respect to particular constraints. Also, more precise orthodontic control over the effect of the aligners can be achieved and their behavior can be better predicted than would otherwise be the case. In addition, computationally defining the aligner geometry facilitates direct aligner manufacturing under numerical control.

The details of one or more embodiments of the invention are set forth in the accompanying drawings and the description below. Other features and advantages of the invention will become apparent from the description, the drawings, and the claims.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1A shows one exemplary dental data mining system.

FIG. 1B shows an analysis of the performance of one or more dental appliances.

FIG. 1C shows various Movement Type data used in one embodiment of the data mining system.

FIG. 1D shows an analysis of the performance of one or more dental appliances.

FIGS. 1E-1F show various embodiments of a clusterizer to generate treatment plans.

FIG. 2A is a flowchart of a process of specifying a course of treatment including a subprocess for calculating aligner shapes in accordance with the invention.

FIG. 2B is a flowchart of a process for calculating aligner shapes.

FIG. 3 is a flowchart of a subprocess for creating finite element models.

FIG. 4 is a flowchart of a subprocess for computing aligner changes.

FIG. 5A is a flowchart of a subprocess for calculating changes in aligner shape.

FIG. 5B is a flowchart of a subprocess for calculating changes in aligner shape.

FIG. 5C is a flowchart of a subprocess for calculating changes in aligner shape.

FIG. 5D is a schematic illustrating the operation of the subprocess of FIG. 5B.

FIG. 6 is a flowchart of a process for computing shapes for sets of aligners.

FIG. 7 is an exemplary diagram of a statistical root model.

FIG. 8 are exemplary diagrams of root modeling.

FIG. 9 are exemplary diagrams of CT scan of teeth.

FIG. 10 shows an exemplary user interface showing teeth.

FIG. 11 shows the exemplary diagram of FIG. 10 with root data.

Like reference numbers and designations in the various drawings indicate like elements.

DETAILED DESCRIPTION OF THE INVENTION

F Digital treatment plans are now possible with 3-dimensional orthodontic treatment planning tools such as ClinCheck® from Align Technology, Inc. or other software available from eModels and OrthoCAD, among others. These technologies allow the clinician to use the actual patient's dentition as a starting point for customizing the treatment plan. The ClinCheck® technology uses a patient-specific digital model to plot a treatment plan, and then use a scan of the achieved treatment outcome to assess the degree of success of the outcome as compared to the original digital treatment plan (previously filed patent for this technology – superimposition tool). The problem with the digital treatment plan and outcome assessment is the abundance of data and the lack of standards and efficient methodology by which to assess "treatment success" at a individual patient level. To analyze the information, a dental data mining system is used.

FIG. 1A shows one exemplary dental data mining system. In this system, dental treatment and outcome data sets 1 are stored in a database or information warehouse 2. The data is extracted by a data mining software 3 that generates results 4. The data mining software can interrogate the information captured and/or updated in the database 2 and can generate an output data stream correlating a patient tooth problem with a dental appliance solution. Note that the output of the data mining software can be most advantageously, self-reflexively, fed as a subsequent input to at least the database and the data mining correlation algorithm.

The result of the data mining system of FIG. 1 is used for defining appliance configurations or changes to appliance configurations for incrementally moving teeth. The tooth movements will be those normally associated with orthodontic treatment,

including translation in all three orthogonal directions relative to a vertical centerline, rotation of the tooth centerline in the two orthodontic directions ("root angulation" and "torque"), as well as rotation about the centerline.

In one embodiment, the data mining system captures the plan, the start position and the final dental position. The system compares the outcome to the plan, using any treatment methodology including removable appliances as well as fixed appliances such as orthodontic brackets and wires, or even other dental treatment such as comparing achieved to plan for orthognathic surgery (may be patents out there because there exists software that compares outcome facial profile to predictive 2-D images), periodontics, restorative, among others.

In one embodiment, a teeth superimposition tool is used to match treatment files of each arch scan. The refinement scan is superimposed over the initial one to arrive at a match based upon tooth anatomy and tooth coordinate system. After teeth in the two arches are matched, the superimposition tool asks for a reference in order to relate the upper arch to the lower arch. When the option "statistical filtering" is selected, the superimposition tool measures the amount of movement for each tooth by first eliminating as reference the ones that move (determined by the difference in position between the current stage and the previous one) more than one standard deviation either above or below the mean of movement of all teeth. The remaining teeth are then selected as reference to measure movement of each tooth.

FIG. 1B shows an analysis of the performance of one or more dental appliances. "Achieved" is plotted against "Goal" in scatter graphs, and trend lines are generated. Scatter graphs are shown to demonstrate where all "scattered" data points are, and trend

lines are generated to show the performance of the dental appliances. In one embodiment, trend lines are selected to be linear (they can be curvilinear); thus trend lines present as the "best fit" straight lines for all "scattered" data. The performance of the Aligners is represented as the slope of a trend line. The Y axis intercept is the incidental movement that occurs when wearing the Aligners. Predictability is measured by $R2$ that is obtained from a regression computation of "Achieved" and "Goal." data. A number of scatter graphs are shown below.

FIG. 1C shows various Movement Type data used in one embodiment of the data mining system. Exemplary data sets cover Expansion/Constriction (+/-X Translation), Mesialization/Distalization (+/-Y Translation), Intrusion (-Z Translation), Extrusion (+Z Translation), Tip/Angulation (X Rotation), Torque/Inclination (Y Rotation), and Pure Rotation (Z Rotation).

FIG. 1D shows an analysis of the performance of one or more dental appliances. FIG. 1D shows that Incisor Intrusions are well controlled, that is, the target goal is achieved about 85% of the time for that particular set of data.

As illustrated saliently in FIG. 1D, actual tooth movement generally lags targeted tooth movement at many stages. In the case of treatment with sequences of polymer appliances, such lags play and important role in treatment design, because both tooth movement and such negative outcomes as patient discomfort vary positively with the extent of the discrepancies.

In one embodiment, clinical parameters in steps such as 170 and 232 are made more precise and safer by allowing for the statistical deviation of targeted from actual tooth position. For example, a subsequent movement target might be reduced because of

a large calculated probability of currently targeted tooth movement not having been achieved adequately, with the result that there is a high probability the subsequent movement stage will need to complete work intended for an earlier stage. Similarly, targeted movement might overshoot desired positions especially in earlier stages so that expected actual movement is better controlled. This embodiment sacrifices the goal of minimizing round trip time in favor of achieving a higher probability of targeted end-stage outcome. This methodology is accomplished within treatment plans specific to clusters of similar patient cases.

Table 1 shows grouping of teeth in one embodiment. The sign convention of tooth movements is indicated in Table 2. Different tooth movements of the selected 60 arches were demonstrated in Table 3 with performance sorted by descending order. The appliance performance can be broken into 4 separate groups: high (79-85%), average (60-68%), below average (52-55%), and inadequate (24-47%). Table 4 shows ranking of predictability. Predictability is broken into 3 groups: highly predictable (.76-.82), predictable (.43-.63) and unpredictable (.10-.30). For the particular set of data, the findings are as follows:

1. Incisor intrusion, and anterior intrusion performance are high. The range for incisor intrusion is about 1.7mm, and for anterior intrusion is about 1.7mm. These movements are highly predictable.

2. Canine intrusion, incisor torque, incisor rotation and anterior torque performance are average. The range for canine intrusion is about 1.3mm, for incisor torque is about 34 degrees, for incisor rotation is about 69 degrees, and for anterior torque is about 34 degrees. These movements are either predictable or highly predictable.

3. Bicuspid tipping, bicuspid mesialization, molar rotation, and posterior expansion performance are below average. The range for bicuspid mesialization is about 1 millimeter, for bicuspid tipping is about 19 degrees, for molar rotation is about 27 degrees and for posterior expansion is about 2.8 millimeters. Bicuspid tipping and mesialization are unpredictable. Whereas the rest are predictable movements.

4. Anterior and incisor extrusion, round teeth and bicuspid rotation, canine tipping, molar distalization, posterior torque performance are inadequate. The range of anterior extrusion is about 1.7 millimeters, for incisor extrusion is about 1.5mm, for round teeth rotation is about 67 degrees for bicuspid rotation is about 63 degrees, for canine tipping is about 26 degrees, for molar distalization is about 2 millimeters, and for posterior torque is about 43 degrees. All are unpredictable movement except bicuspid rotation which is predictable.

Teeth	
Incisors	# 7,8,9,10,23,24,25,26
Canines	# 6,11,22,27
Bicuspids	# 4,5,12,13,20,21,28,29
Molars	# 2,3,14,15,18,19,30,31
Anteriors	# 6,7,8,9,10,11,22,23,24,25,26,27
Posteriors	# 2,3,4,5,12,13,14,15,18,19,20,21,28,29,30,31
Round	# 4,5,6,11,12,13,20,21,22,27,28,29

Table 1.Studied groups of teeth

Type of Movement		
X translation (Expansion/Constriction)	(-) is lingual	(+) is buccal
X rotation (Tipping)		
Upper & Lower right quadrants	(-) is distal	(+) is mesial
Upper & Lower left quadrants	(-) is mesial	(+) is distal
Y translation (Mesialization/Distalization)		
Upper left & Lower right quadrants	(-) is distal	(+) is mesial
Upper right & Lower left quadrants	(-) is mesial	(+) is distal
Y rotation (Torquing)	(-) is lingual crown	(+) is buccal crown
Z translation (Intrusion/Extrusion)	(-) is intrusion	(+) is extrusion
Z rotation (Pure Rotation)	(-) is clockwise	(+) is counterclockwise

Table 2.Sign convention of tooth movements

Attorney Docket No.: AT-000224 US

Group	Movement	Model	Performance Index	Side Effect	Predictability
Incisor	Intrusion	Linear	85%	0.03	0.82
Anterior	Intrusion	Linear	79%	0.03	0.76
Canine	Intrusion	Linear	68%	-0.10	0.43
Incisor	Torque	Linear	67%	0.21	0.63
Anterior	Torque	Linear	62%	0.15	0.56
Incisor	Rotation	Linear	61%	-0.09	0.76
Bicuspid	Tipping	Linear	55%	0.35	0.27
Molar	Rotation	Linear	52%	0.11	0.58
Posterior	Expansion	Linear	52%	0.11	0.48
Bicuspid	Mesialization	Linear	52%	0.00	0.30
Bicuspid	Rotation	Linear	47%	0.28	0.63
Molar	Distalization	Linear	43%	0.02	0.20
Canine	Tipping	Linear	42%	0.10	0.28
Posterior	Torque	Linear	42%	1.50	0.28
Round	Rotation	Linear	39%	-0.14	0.27
Anterior	Extrusion	Linear	29%	-0.02	0.13
Incisor	Extrusion	Linear	24%	0.02	0.10

Table 4.Ranking of Performance Index of movement

Group	Movement	Model	Performance Index	Side Effect	Predictability
Incisor	Intrusion	Linear	85%	0.03	0.82
Anterior	Intrusion	Linear	79%	0.03	0.76
Incisor	Rotation	Linear	61%	-0.09	0.76
Incisor	Torque	Linear	67%	0.21	0.63
Bicuspid	Rotation	Linear	47%	0.28	0.63
Molar	Rotation	Linear	52%	0.11	0.58
Anterior	Torque	Linear	62%	0.15	0.56
Posterior	Expansion	Linear	52%	0.11	0.48
Canine	Intrusion	Linear	68%	-0.10	0.43
Bicuspid	Mesialization	Linear	52%	0.00	0.30
Canine	Tipping	Linear	42%	0.10	0.28
Posterior	Torque	Linear	42%	1.50	0.28
Bicuspid	Tipping	Linear	55%	0.35	0.27
Round	Rotation	Linear	39%	-0.14	0.27
Molar	Distalization	Linear	43%	0.02	0.20
Anterior	Extrusion	Linear	29%	-0.02	0.13
Incisor	Extrusion	Linear	24%	0.02	0.10

Table 4.Ranking of movement predictability

In one embodiment, data driven analyzers may be applied. These data driven analyzers may incorporate a number of models such as parametric statistical models, non-parametric statistical models, clustering models, nearest neighbor models, regression methods, and engineered (artificial) neural networks. Prior to operation, data driven analyzers or models are built using one or more training sessions. The data used to build the analyzer or model in these sessions are typically referred to as training data. As data driven analyzers are developed by examining only training examples, the selection of the training data can significantly affect the accuracy and the learning speed of the data driven analyzer. One approach used heretofore generates a separate data set referred to as a test set for training purposes. The test set is used to avoid overfitting the model or analyzer to the training data. Overfitting refers to the situation where the analyzer has memorized the training data so well that it fails to fit or categorize unseen data. Typically, during the construction of the analyzer or model, the analyzer's performance is tested against the test set. The selection of the analyzer or model parameters is performed iteratively until the performance of the analyzer in classifying the test set reaches an optimal point. At this point, the training process is completed. An alternative to using an independent training and test set is to use a methodology called cross-validation. Cross-validation can be used to determine parameter values for a parametric analyzer or model for a non-parametric analyzer. In cross-validation, a single training data set is selected. Next, a number of different analyzers or models are built by presenting different parts of the training data as test sets to the analyzers in an iterative process. The parameter or model structure is then determined on the basis of the combined performance of all

models or analyzers. Under the cross-validation approach, the analyzer or model is typically retrained with data using the determined optimal model structure.

In one embodiment, the data mining software 3 can be a "spider" or "crawler" to grab data on the database 2 for indexing. In one embodiment, clustering operations are performed to detect patterns in the data. In another embodiment, a neural network is used to recognize each pattern as the neural network is quite robust at recognizing dental treatment patterns. Once the treatment features have been characterized, the neural network then compares the input dental information with stored templates of treatment vocabulary known by the neural network recognizer, among others. The recognition models can include a Hidden Markov Model (HMM), a dynamic programming model, a neural network, a fuzzy logic, or a template matcher, among others. These models may be used singly or in combination.

Dynamic programming considers all possible points within the permitted domain for each value of i. Because the best path from the current point to the next point is independent of what happens beyond that point. Thus, the total cost of [i(k), j(k)] is the cost of the point itself plus the cost of the minimum path to it. Preferably, the values of the predecessors can be kept in an MxN array, and the accumulated cost kept in a 2xN array to contain the accumulated costs of the immediately preceding column and the current column. However, this method requires significant computing resources.

Dynamic programming requires a tremendous amount of computation. For the recognizer to find the optimal time alignment between a sequence of frames and a sequence of node models, it must compare most frames against a plurality of node models. One method of reducing the amount of computation required for dynamic

programming is to use pruning. Pruning terminates the dynamic programming of a given

portion of dental treatment information against a given treatment model if the partial

probability score for that comparison drops below a given threshold. This greatly reduces

computation.

Considered to be a generalization of dynamic programming, a hidden Markov

model is used in the preferred embodiment to evaluate the probability of occurrence of a

sequence of observations O(1), O(2), ... O(t), ..., O(T), where each observation O(t) may

be either a discrete symbol under the VQ approach or a continuous vector. The sequence

of observations may be modeled as a probabilistic function of an underlying Markov

chain having state transitions that are not directly observable.

In the preferred embodiment, the Markov network is used to model a number of

dental treatment options. The transitions between states are represented by a transition

matrix A = [a(i,j)]. Each a(i,j) term of the transition matrix is the probability of making a

transition to state j given that the model is in state i. The output symbol probability of the

model is represented by a set of functions B=[b(j) (O(t)], where the b(j) (O(t) term of the

output symbol matrix is the probability of outputting observation O(t), given that the

model is in state j. The first state is always constrained to be the initial state for the first

time frame of the utterance, as only a prescribed set of left to right state transitions are

possible. A predetermined final state is defined from which transitions to other states

cannot occur.

Transitions are restricted to reentry of a state or entry to one of the next two

states. Such transitions are defined in the model as transition probabilities. For example,

a treatment pattern currently having a frame of feature signals in state 2 has a probability

of reentering state 2 of a(2,2), a probability a(2,3) of entering state 3 and a probability of

a(2,4) =1   a(2,1)   a(2,2) of entering state 4. The probability a(2,1) of entering state 1 or

the probability a(2,5) of entering state 5 is zero and the sum of the probabilities a(2,1)

through a(2,5) is one. Although the preferred embodiment restricts the flow graphs to the

present state or to the next two states, one skilled in the art can build an HMM model

without any transition restrictions, although the sum of all the probabilities of

transitioning from any state must still add up to one.

In each state of the model, the current feature frame may be identified with one of

a set of predefined output symbols or may be labeled probabilistically. In this case, the

output symbol probability b(j) O(t) corresponds to the probability assigned by the model

that the feature frame symbol is O(t). The model arrangement is a matrix A=[a(i,j)] of

transition probabilities and a technique of computing B=b(j) O(t), the feature frame

symbol probability in state j.

The Markov model is formed for a reference pattern from a plurality of sequences

of training patterns and the output symbol probabilities are multivariate Gaussian

function probability densities. The dental treatment information traverses through the

feature extractor. During learning, the resulting feature vector series is processed by a

parameter estimator, whose output is provided to the hidden Markov model. The hidden

Markov model is used to derive a set of reference pattern templates, each template

representative of an identified pattern in a vocabulary set of reference treatment patterns.

The Markov model reference templates are next utilized to classify a sequence of

observations into one of the reference patterns based on the probability of generating the

observations from each Markov model reference pattern template. During recognition,

the unknown pattern can then be identified as the reference pattern with the highest probability in the likelihood calculator.

The HMM template has a number of states, each having a discrete value. However, because treatment pattern features may have a dynamic pattern in contrast to a single value. The addition of a neural network at the front end of the HMM in an embodiment provides the capability of representing states with dynamic values. The input layer of the neural network comprises input neurons. The outputs of the input layer are distributed to all neurons in the middle layer. Similarly, the outputs of the middle layer are distributed to all output states, which normally would be the output layer of the neuron. However, each output has transition probabilities to itself or to the next outputs, thus forming a modified HMM. Each state of the thus formed HMM is capable of responding to a particular dynamic signal, resulting in a more robust HMM. Alternatively, the neural network can be used alone without resorting to the transition probabilities of the HMM architecture.

The output streams or results 4 of FIG. 1A are used as feedback in improving dental appliance design and/or usage by doctors. For example, the data mining results can be used to evaluate performance based on staging approaches, to compare appliance performance indices based on treatment approaches, and to evaluate performance comparing different attachment shapes and positions on teeth.

The ability to study tooth-specific efficacy and product performance for large clusters of treatment outcomes enables statistically significant comparisons to be made between two or more populations of cases. In the event that the two clusters studied contain differences in treatment approach, appliance design, or manufacturing protocol,

the differences seen in the performance of the product as exhibited by the data output, can be attributed to the approach, design, or manufacturing protocol. The end result is a feedback mechanism that enables either the clinician or the manufacturer the ability to optimize the product design and usage based on performance data from a significantly large sample size using objective measurable data.

The theory of orthodontic treatment is not universally agreed upon, and actual treatment and outcomes are subject to additional uncertainties of measurement of patient variables, of relationships to unmeasured patient variables, as well as of varying patient compliance. As a result, different clinicians might prefer different treatment plans for a single patient. Thus, a single treatment plan may not be accepted by every clinician since there is no universally accepted "correct" treatment plan.

The next few embodiments allow greater clinician satisfaction and greater patient satisfaction by tailoring treatment parameters to preferences of clinicians. The system detects differences in treatment preferences by statistical observation of the treatment histories of clinicians. For example, clinicians vary in how likely they would be to perform bicuspid extraction in cases with comparable crowding. Even when there is not a sufficient record of past treatments for a given clinician, clustering may be performed on other predictor variables such as geographical location, variables related to training, or size and nature of practice, to observe statistically significant differences in treatment parameters.

Data mining can discover statistically significant patterns of different treatment outcome achieved by different clinicians for comparable patients. For example, patient cases clustered together might have systematically fewer complications with one clinician

as compared to another. Such a difference detected by the data mining tool might be used as a flag for feedback to the more poorly performing clinician as well as a flag for solicitation of treatment differences used by the better performing clinician.

In one embodiment, clustering techniques are used with previously completed cases to categorize treatment complications and outcomes. Probability models of risk are then built within each cluster. New cases are then allocated to the same clusters based on similarity of pre-treatment variables. The risks within each cluster of patients with completed treatments are then used with new cases to predict treatment outcomes and risks of complications. High-risk patients are then flagged for special attention, possibly including additional steps in treatment plan or additional clinical intervention.

In another embodiment, practitioners are clustered into groups by observed clinician treatment preferences, and treatment parameters are adjusted within each group to coincide more closely with observed treatment preferences. Practitioners without observed histories are then assigned to groups based on similarity of known variables to those within clusters with known treatment histories.

FIG. 1E shows an exemplary process for clusterizing practices. First, the process clusterizes treatment practice based on clinician treatment history such as treatment preferences, outcomes, and demographic & practice variables (20). Next, the system models preferred clinical constraints within each cluster (22). Next, the system assigns clinicians without treatment history to clusters in 20 based on demographic and practice variables (24). In one embodiment, the system performs process 100 separately within each cluster, using cluster-specific clinical constraints (26). Additionally, the system update clusters and cluster assignment as new treatment and outcomes data arrive (28).

Fig. 1F shows another embodiment of a data mining system to generate proposed treatments. First, the system identifies/clusterizes patient histories having detailed follow-up (such as multiple high-resolution scans), based on detailed follow-up data, diagnosis, treatment parameters and outcomes, and demographic variables (40). Within each cluster, the system models discrepancies between intended position and actual positions obtained from follow-up data (42). Further, within each cluster, the system models risk for special undesirable outcomes (44). Patient histories are clusterized with less detailed follow-up data based on available variables. Assign to clusters calculated in 40 (46). The system refines step 42 models with additional records from step 46 clusters (48). It can also refine step 44 models with additional records from step 48 clusters (50). The system then assigns new patients to step 46 clusters based on diagnosis, demographic, and initial physical (52). Within each step 52 cluster, the system models expected discrepancies between intended position and actual positions (54). From step 54, the system uses revised expected position information where relevant (including 232, 250) (67). Additionally, within each step 52 cluster, the system models risk for undesirable outcomes (56). From step 56, the system also flags cases that require special attention and clinical constraints (as in 204, 160) (69). The process then customizes treatment plan to each step 52 cluster (58). Next, the system iteratively collects data (61) and loops back to (40). The system also continually identifies clusters without good representation in step 40 clusters for additional follow-up analysis (65).

In clinical treatment settings, it is not cost-effective to obtain or process the full high-resolution data possible at every stage of tooth movement. For example:

- Patients may use several appliances between visits to clinicians.

- A given patient may submit only one set of tooth impressions.

- Radiation concerns may limit the number of CT or X-Ray scans used.

- Clinicians generally do not have the time to report detailed spatial information on each tooth at each visit.

Due to these and other limitation, treatment planning is necessarily made based on partial information.

In one embodiment, such missing information is approximated substantially by matching predictive characteristics between patients and a representative sample for which detailed follow-up information is collected. In this case, patients are flagged based on poorly anticipated treatment outcomes for requests for follow-up information, such as collection and analysis of additional sets of tooth impressions. Resulting information is then used to refine patient clusters and treatment of patients later assigned to the clusters.

In general, patient data is scanned and the data is analyzed using the data mining system described above. A treatment plan is proposed by the system for the dental practitioner to approve. The dental practitioner can accept or request modifications to the treatment plan. Once the treatment plan is approved, manufacturing of appliance(s) can begin.

FIG. 2A illustrates the general flow of an exemplary process 100 for defining and generating repositioning appliances for orthodontic treatment of a patient. The process 100 includes the methods, and is suitable for the apparatus, of the present invention, as will be described. The computational steps of the process are advantageously implemented as computer program modules for execution on one or more conventional digital computers.

As an initial step, a mold or a scan of patient's teeth or mouth tissue is acquired (110). This step generally involves taking casts of the patient's teeth and gums, and may also involve taking wax bites, direct contact scanning, x-ray imaging, tomographic imaging, sonographic imaging, and other techniques for obtaining information about the position and structure of the teeth, jaws, gums and other orthodontically relevant tissue. From the data so obtained, a digital data set is derived that represents the initial (that is, pretreatment) arrangement of the patient's teeth and other tissues.

The initial digital data set, which may include both raw data from scanning operations and data representing surface models derived from the raw data, is processed to segment the tissue constituents from each other (step 120). In particular, in this step, data structures that digitally represent individual tooth crowns are produced. Advantageously, digital models of entire teeth are produced, including measured or extrapolated hidden surfaces and root structures.

The desired final position of the teeth-- that is, the desired and intended end result of orthodontic treatment -- can be received from a clinician in the form of a prescription, can be calculated from basic orthodontic principles, or can be extrapolated computationally from a clinical prescription (step 130). With a specification of the desired final positions of the teeth and a digital representation of the teeth themselves, the final position and surface geometry of each tooth can be specified (step 140) to form a complete model of the teeth at the desired end of treatment. Generally, in this step, the position of every tooth is specified. The result of this step is a set of digital data structures that represents an orthodontically correct repositioning of the modeled teeth relative to presumed-stable tissue. The teeth and tissue are both represented as digital data.

Having both a beginning position and a final position for each tooth, the process next defines a tooth path for the motion of each tooth. The tooth paths are optimized in the aggregate so that the teeth are moved in the quickest fashion with the least amount of round-tripping to bring the teeth from their initial positions to their desired final positions. (Round-tripping is any motion of a tooth in any direction other than directly toward the desired final position. Round-tripping is sometimes necessary to allow teeth to move past each other.) The tooth paths are segmented. The segments are calculated so that each tooth's motion within a segment stays within threshold limits of linear and rotational translation. In this way, the end points of each path segment can constitute a clinically viable repositioning, and the aggregate of segment end points constitute a clinically viable sequence of tooth positions, so that moving from one point to the next in the sequence does not result in a collision of teeth.

The threshold limits of linear and rotational translation are initialized, in one implementation, with default values based on the nature of the appliance to be used. More individually tailored limit values can be calculated using patient-specific data. The limit values can also be updated based on the result of an appliance-calculation (step 170, described later), which may determine that at one or more points along one or more tooth paths, the forces that can be generated by the appliance on the then-existing configuration of teeth and tissue is incapable of effecting the repositioning that is represented by one or more tooth path segments. With this information, the subprocess defining segmented paths (step 150) can recalculate the paths or the affected subpaths.

At various stages of the process, and in particular after the segmented paths have been defined, the process can, and generally will, interact with a clinician responsible for

the treatment of the patient (step 160). Clinician interaction can be implemented using a client process programmed to receive tooth positions and models, as well as path information from a server computer or process in which other steps of process 100 are implemented. The client process is advantageously programmed to allow the clinician to display an animation of the positions and paths and to allow the clinician to reset the final positions of one or more of the teeth and to specify constraints to be applied to the segmented paths. If the clinician makes any such changes, the subprocess of defining segmented paths (step 150) is performed again.

The segmented tooth paths and associated tooth position data are used to calculate clinically acceptable appliance configurations (or successive changes in appliance configuration) that will move the teeth on the defined treatment path in the steps specified by the path segments (step 170). Each appliance configuration represents a step along the treatment path for the patient. The steps are defined and calculated so that each discrete position can follow by straight-line tooth movement or simple rotation from the tooth positions achieved by the preceding discrete step and so that the amount of repositioning required at each step involves an orthodontically optimal amount of force on the patient's dentition. As with the path definition step, this appliance calculation step can include interactions and even iterative interactions with the clinician (step 160). The operation of a process step 200 implementing this step will be described more fully below.

Having calculated appliance definitions, the process 100 can proceed to the manufacturing step (step 180) in which appliances defined by the process are manufactured, or electronic or printed information is produced that can be used by a

manual or automated process to define appliance configurations or changes to appliance configurations.

FIG. 2B illustrates a process 200 implementing the appliance-calculation step (FIG. 2A, step 170) for polymeric shell aligners of the kind described in above-mentioned patent application no. 09/169,276 (attorney docket no. 018563-004800). Inputs to the process include an initial aligner shape 202, various control parameters 204, and a desired end configuration for the teeth at the end of the current treatment path segment 206. Other inputs include digital models of the teeth in position in the jaw, models of the jaw tissue, and specifications of an initial aligner shape and of the aligner material. Using the input data, the process creates a finite element model of the aligner, teeth and tissue, with the aligner in place on the teeth (step 210). Next, the process applies a finite element analysis to the composite finite element model of aligner, teeth and tissue (step 220). The analysis runs until an exit condition is reached, at which time the process evaluates whether the teeth have reached the desired end position for the current path segment, or a position sufficiently close to the desired end position (step 230). If an acceptable end position is not reached by the teeth, the process calculates a new candidate aligner shape (step 240). If an acceptable end position is reached, the motions of the teeth calculated by the finite elements analysis are evaluated to determine whether they are orthodontically acceptable (step 232). If they are not, the process also proceeds to calculate a new candidate aligner shape (step 240). If the motions are orthodontically acceptable and the teeth have reached an acceptable position, the current aligner shape is compared to the previously calculated aligner shapes. If the current shape is the best solution so far (decision step 250), it is saved as the best candidate so far (step

260). If not, it is saved in an optional step as a possible intermediate result (step 252). If the current aligner shape is the best candidate so far, the process determines whether it is good enough to be accepted (decision step 270). If it is, the process exits. Otherwise, the process continues and calculates another candidate shape (step 240) for analysis.

The finite element models can be created using computer program application software available from a variety of vendors. For creating solid geometry models, computer aided engineering (CAE) or computer aided design (CAD) programs can be used, such as the AutoCAD® software products available from Autodesk, Inc., of San Rafael, Calif. For creating finite element models and analyzing them, program products from a number of vendors can be used, including the PolyFEM product available from CADSI of Coralville, Iowa, the Pro/Mechanica simulation software available from Parametric Technology Corporation of Waltham, Mass., the I-DEAS design software products available from Structural Dynamics Research Corporation (SDRC) of Cincinnati, Ohio, and the MSC/NASTRAN product available from MacNeal-Schwendler Corporation of Los Angeles, Calif.

FIG. 3 shows a process 300 of creating a finite element model that can be used to perform step 210 of the process 200 (FIG. 2). Input to the model creation process 300 includes input data 302 describing the teeth and tissues and input data 304 describing the aligner. The input data describing the teeth 302 include the digital models of the teeth; digital models of rigid tissue structures, if available; shape and viscosity specifications for a highly viscous fluid modeling the substrate tissue in which the teeth are embedded and to which the teeth are connected, in the absence of specific models of those tissues;

and boundary conditions specifying the immovable boundaries of the model elements. In one implementation, the model elements include only models of the teeth, a model of a highly viscous embedding substrate fluid, and boundary conditions that define, in effect, a rigid container in which the modeled fluid is held. Note that fluid characteristics may differ by patient clusters, for example as a function of age.

A finite element model of the initial configuration of the teeth and tissue is created (step 310) and optionally cached for reuse in later iterations of the process (step 320). As was done with the teeth and tissue, a finite element model is created of the polymeric shell aligner (step 330). The input data for this model includes data specifying the material of which the aligner is made and the shape of the aligner (data input 304).

The model aligner is then computationally manipulated to place it over the modeled teeth in the model jaw to create a composite model of an in-place aligner (step 340). Optionally, the forces required to deform the aligner to fit over the teeth, including any hardware attached to the teeth, are computed and used as a figure of merit in measuring the acceptability of the particular aligner configuration. Optionally, the tooth positions used are as estimated from a probabilistic model based on prior treatment steps and other patient information. In a simpler alternative, however, the aligner deformation is modeled by applying enough force to its insides to make it large enough to fit over the teeth, placing the model aligner over the model teeth in the composite model, setting the conditions of the model teeth and tissue to be infinitely rigid, and allowing the model aligner to relax into position over the fixed teeth. The surfaces of the aligner and the teeth are modeled to interact without friction at this stage, so that the aligner model achieves

the correct initial configuration over the model teeth before finite element analysis is begun to find a solution to the composite model and compute the movement of the teeth under the influence of the distorted aligner.

FIG. 4 shows a process 400 for calculating the shape of a next aligner that can be used in the aligner calculations, step 240 of process 200 (FIG. 2). A variety of inputs are used to calculate the next candidate aligner shape. These include inputs 402 of data generated by the finite element analysis solution of the composite model and data 404 defined by the current tooth path. The data 402 derived from the finite element analysis includes the amount of real elapsed time over which the simulated repositioning of the teeth took place; the actual end tooth positions calculated by the analysis; the maximum linear and torsional force applied to each tooth; the maximum linear and angular velocity of each tooth. From the input path information, the input data 404 includes the initial tooth positions for the current path segment, the desired tooth positions at the end of the current path segment, the maximum allowable displacement velocity for each tooth, and the maximum allowable force of each kind for each tooth.

If a previously evaluated aligner was found to violate one or more constraints, additional input data 406 can optionally be used by the process 400. This data 406 can include information identifying the constraints violated by, and any identified suboptimal performance of, the previously evaluated aligner.

Having received the initial input data (step 420), the process iterates over the movable teeth in the model. (Some of the teeth may be identified as, and constrained to be, immobile.) If the end position and dynamics of motion of the currently selected tooth

by the previously selected aligner is acceptable ("yes" branch of decision step 440), the process continues by selecting for consideration a next tooth (step 430) until all teeth have been considered ("done" branch from step 430 to step 470). Otherwise ("no" branch from step 440), a change in the aligner is calculated in the region of the currently selected tooth (step 450). The process then moves back to select the next current tooth (step 430) as has been described.

When all of the teeth have been considered, the aggregate changes made to the aligner are evaluated against previously defined constraints (step 470), examples of which have already been mentioned. Constraints can be defined with reference to a variety of further considerations, such as manufacturability. For example, constraints can be defined to set a maximum or minimum thickness of the aligner material, or to set a maximum or minimum coverage of the aligner over the crowns of the teeth. If the aligner constraints are satisfied, the changes are applied to define a new aligner shape (step 490). Otherwise, the changes to the aligner are revised to satisfy the constraints (step 480), and the revised changes are applied to define the new aligner shape (step 490).

FIG. 5A illustrates one implementation of the step of computing an aligner change in a region of a current tooth (step 450). In this implementation, a rule-based inference engine 456 is used to process the input data previously described (input 454) and a set of rules 452a-452n in a rule base of rules 452. The inference engine 456 and the rules 452 define a production system which, when applied to the factual input data, produces a set of output conclusions that specify the changes to be made to the aligner in the region of the current tooth (output 458).

Rules 452 have the conventional two-part form: an if-part defining a condition and a then-part defining a conclusion or action that is asserted if the condition is satisfied. Conditions can be simple or they can be complex conjunctions or disjunctions of multiple assertions. An exemplary set of rules, which defines changes to be made to the aligner, includes the following: if the motion of the tooth is too slow, add driving material to the aligner opposite the desired direction of motion; if the motion of the tooth is too slow, add driving material to overcorrect the position of the tooth; if the tooth is too far short of the desired end position, add material to overcorrect; if the tooth has been moved too far past the desired end position, add material to stiffen the aligner where the tooth moves to meet it; if a maximum amount of driving material has been added, add material to overcorrect the repositioning of the tooth and do not add driving material; if the motion of the tooth is in a direction other than the desired direction, remove and add material so as to redirect the tooth.

In an alternative embodiment, illustrated in FIGS. 5B and 5C, an absolute configuration of the aligner is computed, rather than an incremental difference. As shown in FIG. 5B, a process 460 computes an absolute configuration for an aligner in a region of a current tooth. Using input data that has already been described, the process computes the difference between the desired end position and the achieved end position of the current tooth (462). Using the intersection of the tooth center line with the level of the gum tissue as the point of reference, the process computes the complement of the difference in all six degrees of freedom of motion, namely three degrees of translation and three degrees of rotation (step 464). Next, the model tooth is displaced from its

desired end position by the amounts of the complement differences (step 466), which is illustrated in FIG. 5D.

FIG. 5D shows a planar view of an illustrative model aligner 60 over an illustrative model tooth 62. The tooth is in its desired end position and the aligner shape is defined by the tooth in this end position. The actual motion of the tooth calculated by the finite element analysis is illustrated as placing the tooth in position 64 rather than in the desired position 62. A complement of the computed end position is illustrated as position 66. The next step of process 460 (FIG. 5B) defines the aligner in the region of the current tooth in this iteration of the process by the position of the displaced model tooth (step 468) calculated in the preceding step (466). This computed aligner configuration in the region of the current tooth is illustrated in FIG. 5D as shape 68 which is defined by the repositioned model tooth in position 66.

A further step in process 460, which can also be implemented as a rule 452 (FIG. 5A), is shown in FIG. 5C. To move the current tooth in the direction of its central axis, the size of the model tooth defining that region of the aligner, or the amount of room allowed in the aligner for the tooth, is made smaller in the area away from which the process has decided to move the tooth (step 465).

As shown in FIG. 6, the process 200 of computing the shape for an aligner for a step in a treatment path is one step in an overall process 600 of computing the shapes of a series of aligners. This overall process 600 begins with an initialization step 602 in which initial data, control and constraint values are obtained.

When an aligner configuration has been found for each step or segment of the treatment path (step 604), the overall process 600 determines whether all of the aligners

are acceptable (step 606). If they are, the process exits and is complete. Otherwise, the process optionally undertakes a set of steps 610 in an attempt to calculate a set of acceptable aligners. First, one or more of the constraints on the aligners is relaxed (step 612). Then, for each path segment with an unacceptable aligner, the process 200 of shaping an aligner is performed with the new constraints (step 614). If all the aligners are now acceptable, the overall process 600 exits (step 616).

Aligners may be unacceptable for a variety of reasons, some of which are handled by the overall process. For example, if any impossible movements were required (decision step 620), that is, if the shape calculation process 200 was required to effect a motion for which no rule or adjustment was available, the process 600 proceeds to execute a module that calculates the configuration of a hardware attachment to the subject tooth to which forces can be applied to effect the required motion (step 640). Because adding hardware can have an effect that is more than local, when hardware is added to the model, the outer loop of the overall process 600 is executed again (step 642).

If no impossible movements were required ("no" branch from step 620), the process transfers control to a path definition process (such as step 150, FIG. 2A) to redefine those parts of the treatment path having unacceptable aligners (step 630). This step can include both changing the increments of tooth motion, i.e., changing the segmentation, on the treatment path, changing the path followed by one or more teeth in the treatment path, or both. After the treatment path has been redefined, the outer loop of the overall process is executed again (step 632). The recalculation is advantageously limited to recalculating only those aligners on the redefined portions of the treatment

path. If all the aligners are now acceptable, the overall process exits (step 634). If unacceptable aligners still remain, the overall process can be repeated until an acceptable set of aligners is found or an iteration limit is exceeded (step 650). At this point, as well as at other point in the processes that are described in this specification, such as at the computation of additional hardware (step 640), the process can interact with a human operator, such as a clinician or technician, to request assistance (step 652). Assistance that an operator provides can include defining or selecting suitable attachments to be attached to a tooth or a bone, defining an added elastic element to provide a needed force for one or more segments of the treatment path, suggesting an alteration to the treatment path, either in the motion path of a tooth or in the segmentation of the treatment path, and approving a deviation from or relaxation of an operative constraint.

As was mentioned above, the overall process 600 is defined and parameterized by various items of input data (step 602). In one implementation, this initializing and defining data includes the following items: an iteration limit for the outer loop of the overall process; specification of figures of merit that are calculated to determine whether an aligner is good enough (see FIG. 2, step 270); a specification of the aligner material; a specification of the constraints that the shape or configuration of an aligner must satisfy to be acceptable; a specification of the forces and positioning motions and velocities that are orthodontically acceptable; an initial treatment path, which includes the motion path for each tooth and a segmentation of the treatment path into segments, each segment to be accomplished by one aligner; a specification of the shapes and positions of any anchors installed on the teeth or otherwise; and a specification of a model for the jaw bone and other tissues in or on which the teeth are situated (in the implementation being described,

this model consists of a model of a viscous substrate fluid in which the teeth are embedded and which has boundary conditions that essentially define a container for the fluid).

FIG. 7 is an exemplary diagram of a statistical root model. As shown therein, using the scanning processes described above, a scanned upper portion 700 of a tooth is identified. The scanned upper portion, including the crown, is then supplemented with a modeled 3D root. The 3D model of the root can be statistically modeled. The 3D model of the root 702 and the 3D model of the upper portion 700 together form a complete 3D model of a tooth.

FIG. 8 are exemplary diagrams of root modeling, as enhanced using additional dental information. In FIG. 8, the additional dental information is X-ray information. An X-ray image 710 of teeth is scanned to provide a 2D view of the complete tooth shapes. An outline of a target tooth is identified in the X-Ray image. The model 712 as developed in FIG. 7 is modified in accordance with the additional information. In one embodiment, the tooth model of FIG. 7 is morphed to form a new model 714 that conforms with the X-ray data.

FIG. 9 are exemplary diagrams of a CT scan of teeth. In this embodiment, the roots are derived directly from a high-resolution CBCT scan of the patient. Scanned roots can then be applied to crowns derived from an impression, or used with the existing crowns extracted from Cone Beam Computed Tomography (CBCT) data. A CBCT single scan gives 3D data and multiple forms of X-ray-like data. PVS impressions are avoided.

In one embodiment, a cone beam x-ray source and a 2D area detector scans the patient's dental anatomy, preferably over a 360 degree angular range and along its entire

length, by any one of various methods wherein the position of the area detector is fixed relative to the source, and relative rotational and translational movement between the source and object provides the scanning (irradiation of the object by radiation energy). As a result of the relative movement of the cone beam source to a plurality of source positions (i.e., "views") along the scan path, the detector acquires a corresponding plurality of sequential sets of cone beam projection data (also referred to herein as cone beam data or projection data), each set of cone beam data being representative of x-ray attenuation caused by the object at a respective one of the source positions.

FIG. 10 shows an exemplary user interface showing the erupted teeth, while FIG. 11 shows the exemplary diagram of the teeth of FIG. 10 with root information. Each tooth is individually adjustable using a suitable handle. In the embodiment of FIGS. 10 and 11, the handle allows an operator to move the tooth in three-dimensions with six degrees of freedom.

The teeth movement is guided in part using a root-based sequencing system. In one embodiment, the movement is constrained by a surface area constraint, while in another embodiment, the movement is constrained by a volume constraint.

In one embodiment, the system determines a surface area for each tooth model. The system then sums all surface areas for all tooth models to be moved. Next, the system sums all surface areas of all tooth models on the arch. For each stage of teeth movement, the system checks that a predetermined area ratio or constraint is met while the tooth models are moved. In one implementation, the constraint can be to ensure that the surface areas of moving teeth are less than the total surface areas of teeth on an arch supporting the teeth being moved. If the ratio is greater than a particular number such as

50%, the system indicates an error signal to an operator to indicate that the teeth should be moved on a slower basis.

In another embodiment, the system determines the volume for each tooth model. The system then sums the volumes for all tooth models being moved. Next, the system determines the total volume of all tooth models on the arch. For each stage of teeth movement, the system checks that a predetermined volume ratio or constraint is met while the tooth models are moved. In one implementation, the constraint can be to ensure that the volume for moving teeth is less than the volume of all teeth on an arch supporting the teeth being moved. If the ratio is greater than a particular number such as 50%, the system indicates an error signal to an operator to indicate that the teeth should be moved on a slower basis.

Optionally, other features are added to the tooth model data sets to produce desired features in the aligners. For example, it may be desirable to add digital wax patches to define cavities or recesses to maintain a space between the aligner and particular regions of the teeth or jaw. It may also be desirable to add digital wax patches to define corrugated or other structural forms to create regions having particular stiffness or other structural properties. In manufacturing processes that rely on generation of positive models to produce the repositioning appliance, adding a wax patch to the digital model will generate a positive mold that has the same added wax patch geometry. This can be done globally in defining the base shape of the aligners or in the calculation of particular aligner shapes. One feature that can be added is a rim around the gumline, which can be produced by adding a digital model wire at the gumline of the digital model teeth from which the aligner is manufactured. When an aligner is manufactured by

pressure fitting polymeric material over a positive physical model of the digital teeth, the wire along the gumlines causes the aligner to have a rim around it providing additional stiffness along the gumline.

In another optional manufacturing technique, two sheets of material are pressure fit over the positive tooth model, where one of the sheets is cut along the apex arch of the aligner and the other is overlaid on top. This provides a double thickness of aligner material along the vertical walls of the teeth.

The changes that can be made to the design of an aligner are constrained by the manufacturing technique that will be used to produce it. For example, if the aligner will be made by pressure fitting a polymeric sheet over a positive model, the thickness of the aligner is determined by the thickness of the sheet. As a consequence, the system will generally adjust the performance of the aligner by changing the orientation of the model teeth, the sizes of parts of the model teeth, the position and selection of attachments, and the addition or removal of material (e.g., adding wires or creating dimples) to change the structure of the aligner. The system can optionally adjust the aligner by specifying that one or more of the aligners are to be made of a sheet of a thickness other than the standard one, to provide more or less force to the teeth. On the other hand, if the aligner will be made by a stereo lithography process, the thickness of the aligner can be varied locally, and structural features such as rims, dimples, and corrugations can be added without modifying the digital model of the teeth.

The system can also be used to model the effects of more traditional appliances such as retainers and braces and therefore be used to generate optimal designs and treatment programs for particular patients.

The data processing aspects of the invention can be implemented in digital electronic circuitry, or in computer hardware, firmware, software, or in combinations of them. Data processing apparatus of the invention can be implemented in a computer program product tangibly embodied in a machine-readable storage device for execution by a programmable processor; and data processing method steps of the invention can be performed by a programmable processor executing a program of instructions to perform functions of the invention by operating on input data and generating output. The data processing aspects of the invention can be implemented advantageously in one or more computer programs that are executable on a programmable system including at least one programmable processor coupled to receive data and instructions from and to transmit data and instructions to a data storage system, at least one input device, and at least one output device. Each computer program can be implemented in a high-level procedural or objectoriented programming language, or in assembly or machine language, if desired; and, in any case, the language can be a compiled or interpreted language. Suitable processors include, by way of example, both general and special purpose microprocessors. Generally, a processor will receive instructions and data from a read-only memory and/or a random access memory. Storage devices suitable for tangibly embodying computer program instructions and data include all forms of nonvolatile memory, including by way of example semiconductor memory devices, such as EPROM, EEPROM, and flash memory devices; magnetic disks such as internal hard disks and removable disks; magneto-optical disks; and CD-ROM disks. Any of the foregoing can be supplemented by, or incorporated in, ASICs (application-specific integrated circuits).

To provide for interaction with a user, the invention can be implemented using a computer system having a display device such as a monitor or LCD (liquid crystal display) screen for displaying information to the user and input devices by which the user can provide input to the computer system such as a keyboard, a two-dimensional pointing device such as a mouse or a trackball, or a three-dimensional pointing device such as a data glove or a gyroscopic mouse. The computer system can be programmed to provide a graphical user interface through which computer programs interact with users. The computer system can be programmed to provide a virtual reality, three-dimensional display interface.

The invention has been described in terms of particular embodiments. Other embodiments are within the scope of the following claims. For example, the steps of the invention can be performed in a different order and still achieve desirable results.